120

PATENT AF|2662 #14

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant(s): Hans Eberle and Neil C. Wilhelm

Title: SWITCHED NETWORK FOR LOW LATENCY COMMUNICATION

Application No.: 09/540,730 Filed: March 31, 2000

Examiner: Timothy L. Lee Group Art Unit: 2662

Atty. Docket No.: 004-4255

August 11, 2004

Mail Stop Appeal Brief - Patents
COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, VA 22313-1450

APPELLANT'S BRIEF (37 C.F.R. § 1.192)

This brief is in furtherance of the Notice of Appeal, filed on May 17, 2004. The fees required under § 1.17(c), are provided in the accompanying Transmittal. This brief is being transmitted in triplicate pursuant to 37 C.F.R. § 1.192(a). A Petition for Extension of Time (for a one month extension) is enclosed herewith, extending the response time to August 17, 2004.

REAL PARTY IN INTEREST

The real party in interest in this appeal is Sun Microsystems, Inc., as evidenced by the assignment recorded at Reel 010674/Frame 0597.

RELATED APPEALS AND INTERFERENCES

Appellants have no knowledge of any related appeals or interferences.

STATUS OF CLAIMS

Claims 1-30 are presented herein on appeal. Claims 1-30 were rejected in a final Office action dated January 8, 2004. That final rejection is now appealed.

Claims 1-30 presented herein on appeal are reproduced in the Appendix attached hereto. Claims 1-19 and claims 21-29 were originally presented. Claims 20 and 30 were amended in response to a first Office Action.

STATUS OF AMENDMENTS

No amendments have been filed subsequent to final rejection.

SUMMARY OF INVENTION

In one embodiment, the invention provides a method for communicating packets between one of a plurality of sending nodes and one of a plurality of receiving nodes of a switched network. The switched network includes a buffer-less switch coupling the sending nodes and the receiving nodes. The method includes transmitting packets from respective sending nodes to respective input ports of the buffer-less switch and forwarding all packets that are successfully delivered to the receiving nodes through output ports of the buffer-less switch, through the buffer-less switch with a fixed forwarding rate. The switched network resolves conflicts in requests for a transmission path on the switch from multiple packets by allocating the transmission path to a first requester for the transmission path. If multiple requests for a switch resource collide by requesting the switch resource at the same time, one of the requests is selected as a winner and the one or more remaining packets associated with the losing request(s) are dropped. The winning packet may be selected on a random basis or a round robin basis or based on some other criteria. The requests for transmission paths may be contained within the packets sent into the network and extracted on entry of the packet into the network. A sending node can check for successful transmission of a packet at a predetermined time because a receiving node sends an acknowledge back after a fixed delay. A time out can also be used to determine if a packet was transmitted successfully by checking after the predetermined time for the acknowledge.

In still another embodiment of the invention a computing system is provided that includes a plurality of sending and receiving nodes, each of the sending and receiving nodes including a processor. The computing system also includes a low latency switched network including a buffer-less switch coupling the plurality of sending and receiving nodes. The buffer-less switch

has a fixed forwarding delay for all packets sent from one of the sending nodes and successfully received by one of the receiving nodes because no buffering takes place within the switch in case of contention for resources. The switch includes switch control logic, coupled to input registers in the switch. The switch control logic is responsive to packet information stored in the registers, to allocate output ports on the switch according the packet information on a first come first served basis. Respective packet information provided to the switch control logic constitutes respective requests for output ports, and if a first and second request for an output port path collide by requesting the output port at the same time, the switch control logic responds by selecting one of the requests as a winner and dropping the other packet.

One type of switch appropriate for an embodiment is shown in Figure 3. Referring to Figure 3, a block diagram shows a crossbar switch **300**. Each of input ports **310** is coupled to each of output ports **320**. Assuming each input port **310** and each output port **320** have the same bandwidth "b," resource conflicts can arise. According to an embodiment, if no buffer memory is present in the switch **300** to temporarily store data packets, and multiple data packets are simultaneously forwarded to one of output ports **320**, switch **300** drops data packets.

One method of preventing conflicts requires an input buffer memory or output buffer memory to temporarily store packets. An input buffer holds a data packet in an input buffer coupled to the switch **300** and prevents the data packet from entering the switch **300** until a desired one of the output ports **320** is available. Similarly, output buffering avoids conflicts by providing an output buffer memory with enough input bandwidth to allow packets to be received simultaneously from all input ports **310**.

Referring now to Figure 4, a switched data network includes bulk channel switch **450**, which is a non-blocking buffer-less switch. Note that each node may include separate buffers or queues for the different nodes. In fact, each node may include separate send and/or receive queues for each node on the switch. Referring now to Figure 5, a block diagram shows a non-blocking buffer-less switch **500** that is appropriate for implementing bulk channel switch **450**. The nodes of switched data network **500** queue the data packets outside the buffer-less switch **500**. The term "buffer-less" refers to the fact that the switch provides no buffers for temporarily storing packets or portions of packets in case there are conflicts during a transfer for a particular

switch resource, typically an output port. To avoid conflicts, non-blocking buffer-less switch **500** includes a switch scheduler **510** that controls the scheduling of packets to and from each of network nodes **520, 530, 540** and **550**.

Main memories within the nodes may provide buffering for data packets. Thus, network node **520** includes receive buffer **570** and transmit buffer **560** within a computer system memory. The computer system memory is coupled to a network interface within the computer system that stores a portion of the transmit and receive buffers. The network interface within each node receives commands from switch scheduler **510** governing when to send data packets.

According to another embodiment, each network node **520, 530, 540**, and **550** includes multiple storage queues. Thus, each network node includes a queue for sending packets and a queue for receiving packets, or, alternatively, one or more send queues and receive queues. Thus, each input port couples to a queue and each output port couples to a queue. Each queue disposed within each network node may include a portion of the queue within a network interface.

The switched data network illustrated in Figure 5 requests permission to transmit a packet through a buffer-less switch **500**. More specifically, the request for permission includes communicating with switch scheduler **510** via signal REQ **580**. In response, switch scheduler **510** provides one of a grant or a denial of permission via signal GNT **590**. An assigned transmission slot is received from the switch scheduler **510** via GNT **590**.

The data packet is transferred through the buffer-less switch in an assigned transmission slot. Because there are no buffers in the switch to resolve conflicts, forwarding delays through the switch are fixed. That is, it takes a fixed amount of time for a packet to cross the switch. Being buffer-less does not imply that there can be no storage elements in the switch, it simply means that any switch storage elements that are present do not provide buffering resulting in variable transmission delays through the switch. Thus, any time a portion of a packet is stored in the switch, it is stored for a fixed amount of time before it is forwarded on. That simplifies scheduling of the switch.

Low latency is achieved, in part, by allowing a network to lose packets. Rather than coordinating and scheduling accesses to shared resources, such as registers, buffers, and, in particular, transmission paths, resources are assumed to be always available. In the event of a conflict, one packet wins and the other ones fail. If transmission fails, it is the sender's responsibility to resend the packet. The lossy network scheme avoids time-consuming scheduling operations as long as the network resources are only lightly loaded and conflicts occur infrequently.

A lossy network is particularly attractive since it allows one to build simple and fast switches such as the switch illustrated in Figure 8. Although a 2X2 switch is illustrated for ease of understanding, the concepts described herein associated with a lossy switch can be incorporated into any size switch. No time-consuming arbitration or scheduling of its data paths is required. Packets are forwarded on a first come first served basis. Thus, as shown in Figure 9A, packet B is dropped because it arrived at the output port selector circuit later than packet A. If packets do happen to collide, one packet wins and the other packet(s) are dropped. Thus, as shown in Fig. 9B, packet A is chosen as the loser based on some simple algorithm such as a random or round robin selection. More sophisticated algorithms can be chosen such as selecting the winner according to a fairness criteria having an objective to allocate the same amount of output port bandwidth to each input port on the switch. Any approach used to choose a winner should preferably not add any more than necessary to latency.

Lossy communication also makes it possible to use simple and fast buffering schemes in the sending and receiving nodes of the network. Referring to Figure 10, assume that the sender and the receiver are either a user program, a systems program, or a transmission protocol. Figure 10 again illustrates a buffer-free 2X2 switch **1010**. Assume that node 0 is sending a packet. To send a packet, a node 0 writes a packet into send register **1012**. Node 0 then polls a status register **1014** until it becomes valid. Once the status register is valid, it will indicate whether the transmission was successful. If the status register indicates that the transmission was unsuccessful, the sender has to resend the packet by writing the packet into send register **1012**. Because low latency communication is typically synchronous in that a sender cannot proceed until it is known that the transmission was successful, the sender can be put in charge of doing the retransmission if necessary. Successful and unsuccessful transmission can be determined

with the help of an acknowledge packet (ack) or no acknowledge packet (nack), respectively, or a timeout mechanism in which the sending node waits a predetermined amount of time to see if an acknowledge indicating a successful transmission is received. If not, the sender assumes an error. When the target is node 1, the status register **1014** may receive an ack written into the node 1 send register when node 1 successfully receives the sent packet or may receive a nack when node 1 detects an error in receipt of a packet. The status register is thus coupled to receive information such as an acknowledge or no acknowledge packet received into the node 0 receive buffers. Latency is reduced in that no complicated data structure such as a list of buffers has to be processed.

The unreliable behavior of the network simplifies other parts of the implementation of the network. In one simple implementation, the receiving node drops a packet when it detects a transmission error or when a receive buffer overflows. The transmission error may be detected using, e.g., a checksum or CRC. A timeout mechanism can inform the sender accordingly. A more sophisticated approach reports errors to the sender to allow the system to better determine the cause of packet loss. In any case, the network does not have to be able to retransmit erroneously transmitted packets, as that task is left to the sender.

A further simplification can be achieved by having the receiver send an acknowledge or a nack at a fixed time relative to when the packet is sent. In that way, after a predetermined delay, a sender can check and determine conclusively whether transmission was successful. Either an acknowledge or a nack will be received within the predetermined time period or the sender can conclude that the transfer failed since an acknowledge (or nack) was not received after the fixed delay. Note that in some implementations, a timeout can be used instead of or in addition to a nack. In systems with variable forwarding delays, timeout mechanisms are less reliable as an indication of a transmission failure.

No intermediate buffers are needed between the sender and the receiver, as are typically found in other switching networks. If conflicts occur, rather than buffering a packet, packets are simply dropped. As a consequence, no buffering or buffer management including flow control is needed.

ISSUES

The issues on appeal are:

1. whether claims 13, and 21-28 are unpatentable under 35 U.S.C. § 103(a) over U.S. Patent No. 6,115,373 to Lea in view of U.S. Patent No. 6,072,772 to Charney et al.; and

2. whether claims 1-3, 6-11, 18, 19, and 29 are unpatentable under 35 U.S.C. § 102(e) over U.S. Patent No. 6,115,373 to Lea;

3. whether claims 4, 5, and 12 are unpatentable under 35 U.S.C. § 103(a) over Lea in view of U.S. Patent No. 5,359,320 to Jaffe et al.

4. whether claims 14-17 and 30 are unpatentable under 35 U.S.C. § 103(a) over Lea.

GROUPING OF CLAIMS

Appealed claims do not stand or fall together. The claims are grouped as follows:

Claims Group I: Claims 13, 22-24

Claims Group II: Claims 11 and 12

Claims Group III: Claims 30, 14-17

Claims Group IV: Claims 1-10

Claims Group V: Claims 18-21, 25-29

ARGUMENTS

Claims Group I:

Referring to the first group of claims, these claims have been rejected under 35 U.S.C. § 103(a) as being unpatentable over U.S. Patent No. 6,115,373 to Lea in view of U.S. Patent No. 6.072,772 to Charney et al. For purposes of this discussion, claim 13 of Group I will be addressed. The Office Actions fail to point out where Lea and Charney teach or suggest, and

Appellants respectfully maintain that neither Lea nor Charney, alone or in combination with other references of record, teaches or suggests a sending node polling a status register to determine if transmission of a packet was successful, as recited by claim 13.

The Office Action mailed July 7, 2003 states that intermediary input buffers of Charney "could act as the status register." However, that Office Action and the Office Action mailed January 8, 2004 (hereinafter, the "Final Office Action) fail to provide a reference providing a teaching or suggestion that supports this statement. The input buffers of Charney store packets. (Col. 6, lines 12-34) Nowhere does Charney teach or suggest that these buffers contain status information indicating that the transmission of the packet was successful, or that these buffers are polled to determine if transmission of the packet was successful, as required by claim 13. The Final Office Action states that "[t]he combination of the two references suggests that the acknowledgements could be stored in these registers. It would have been obvious to check or poll these registers to see if there was a successful termination, or there would not point [sic] in keeping these registers and it would be a waste of memory resources." Applicants respectfully disagree. The Office Action fails to provide a reference that teaches or suggests how the input buffers of Charney can be modified to contain status information without rendering the input buffers of Charney unsatisfactory for the intended purpose of queuing input packets. See MPEP §2143.01.

In addition, Lea, alone or in combination with Charney, fails to teach or suggest rewriting the packet to the send register if transmission was unsuccessful. Thus, Lea and Charney fail to teach or suggest, alone or in combination, the claimed invention. For at least these reasons, the PTO's rejection of claim 13 and all claims in Group I should be reversed.

Claims Group II:

Referring to the second group of claims, claim 11 has been rejected under 35 U.S.C. § 102(e) as being anticipated by U.S. Patent No. 6,115,373 to Lea and claim 12 has been rejected under 35 U.S.C. § 103 as being unpatentable over Lea in view of U.S. Patent No. 5,359,320 to Jaffe et al.. With reference to claim 11, Lea fails to teach or suggest that no buffer space is allocated in a receiving node before a packet is sent, thereby simplifying switch overhead, as recited in claim 11. The Office Action states that "Lea does not disclose anything about pre-

allocating buffer space in the buffer before the packet is sent." However, "[a] claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference. *Verdegaal Bros. v. Union Oil Co. of California*, 814 F.2d 628, 631, 2 USPQ2d 1051, 1053 (Fed. Cir. 1987)." See MPEP § 2131.

The Final Office Action refers to Figure 10 and states that "it appears that there aren't even any buffers after the switch matrix, so no system would allocated system [sic] in buffers that don't exist." However, claim 11 recites that no buffer space is allocated in a receiving node. Figure 10 refers to "the switch in more detail" (col. 3, line 27) and is not directed to the receiving node of the system of Lea. Although Figure 10 and col. 5, lines 17-30 show buffers prior to the switch matrix, these portions of Lea fail to teach or suggest, expressly or inherently, that no buffer space is allocated in a receiving node before a packet is sent, thereby simplifying switch overhead as recited in claim 11. For at least this reason, Applicants believe claim 11 is allowable over Lea.

Claim 12 should not be rejected under 35 U.S.C. § 103 as being unpatentable over Lea in view of U.S. Patent No. 5,359,320 to Jaffe et al. because Lea fails to teach the limitations of claim 11 and Jaffe fails to compensate for the shortcomings of Lea. Jaffe teaches a scheduling mechanism for controlling when the arbitration circuit of a node sharing a CSMA communication medium is to start CSMA arbitration for access to the communication medium once the node has a message ready for transmission. (Abstract) Jaffe fails to teach or suggest that no buffer space is allocated in a receiving node before a packet is sent, thereby simplifying switch overhead, as recited in claim 11.

Thus, the references fail to teach or suggest, alone or in combination, the claimed invention. For at least these reasons, the PTO's rejection of claims 11 and 12 should be reversed.

Claims Group III:

Referring to the third group of claims (claims 30 and 14-17), these claims have been rejected under 35 U.S.C. § 103(a) as being unpatentable over Lea. For purposes of this discussion, claim 30 will be addressed. Applicants respectfully maintain that Lea, alone or in combination with other references of record, fails to teach or suggest means for forwarding

packets at a fixed rate on a first come first served basis from respective input ports through a switch to respective output ports, as recited by claim 30. The Office Action states, with reference to claims 1 and 18, that Lea "mentions nothing about changing the forwarding rates of the incoming packets." The Office Action further states that

> Lea discloses in col. 4, lines 1-15, that the system can handle real-time traffic. If the incoming packets are real-time traffic, then they will be forwarded at this same fixed rate as well. Therefore Lea discloses a system that can have a fixed forwarding rate.

Applicants respectfully disagree. A "real-time system" is a "system in which transactions are processed as they occur." RUDOLPH F. GRAF, MODERN DICTIONARY OF ELECTRONICS, at 628 (7th ed. 1999).

The Office Action fails to point out how a system that processes data as it is received teaches or suggests that the switch of Lea has a fixed forwarding rate. "A claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference. *Verdegaal Bros. v. Union Oil Co. of California*, 814 F.2d 628, 631, 2 USPQ2d 1051, 1053 (Fed. Cir. 1987)." See MPEP § 2131 (emphasis added). Lea alone or in combination with other references of record fails to teach or suggest, expressly or inherently, means for forwarding packets at a fixed rate on a first come first served basis from respective input ports through a switch to respective output ports as recited in claim 30.

In addition, Lea fails to teach or suggest forwarding packets on "a first come first served basis" as recited in claim 30. The Final Office Action states that

> Lea does not expressly disclose forwarding packets on a first come first served basis; however, it is well-known in the art that ATM supports constant bit rate services. It would have been obvious to a person of ordinary skill in the art to have the system of Lea support CBR. One would have been motivated to do this because Lea mentions that there are four different classes of services, and one of those classes include real-time traffic, and real-time traffic can depend on constant bit rate to transport data.

The Final Office Action fails to point out how constant bit rate services teach or suggest forwarding packets on a first come first served basis. Lea teaches prioritization of packets at col.

3, line 58-col. 4, line 20. These prioritizations fail to teach that packets are forwarded on a first come first served basis.

For at least this reason, Applicants believe claim 30 is allowable over Lea, alone or in combination with other references of record. Accordingly, Applicants respectfully request that the rejection of independent claim 30, and all claims of Group III, be withdrawn.

Claims Group IV

Referring to the fourth group of claims (claims 1-10), these claims have been rejected under 35 U.S.C. § 102(e) as being unpatentable over U.S. Patent No. 6,115,373 to Lea. For purposes of this discussion, claim 1 of Group IV will be addressed. Applicants respectfully maintain that Lea fails to teach or suggest forwarding all packets that are successfully delivered through output ports of the buffer-less switch to the receiving nodes, through the buffer-less switch with a fixed forwarding rate as recited by claim 1. The Final Office Action states that Lea "mentions nothing about changing the forwarding rates of the incoming packets, so the system forwards the incoming packets at the same rate that they were received." The Final Office Action states that

> if the system of Lea receives packets at a constant bit rate, then the system will be forwarding the incoming packets at the same constant bit rate, which is also known as a fixed rate. Lea discloses in col. 4, lines 1-15, that the system can handle real-time traffic. If the incoming packets are real-time traffic, then they will be forwarded at this same fixed rate as well. Therefore Lea discloses a system that can have a fixed forwarding rate.

Applicants respectfully disagree. A "real-time system" is a "system in which transactions are processed as they occur." RUDOLPH F. GRAF, MODERN DICTIONARY OF ELECTRONICS, at 628 (7th ed. 1999).

The Final Office Action fails to point out how a system that processes data as it is received teaches or suggests that the switch of Lea has a fixed forwarding rate. "A claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference. *Verdegaal Bros. v. Union Oil Co. of California*, 814 F.2d 628, 631, 2 USPQ2d 1051, 1053 (Fed. Cir. 1987)." See MPEP § 2131 (emphasis added). Lea fails to teach or suggest, expressly or inherently, forwarding all packets

that are successfully delivered through output ports of the buffer-less switch to the receiving nodes, through the buffer-less switch with a fixed forwarding rate, as recited by claim 1.

Thus, the references fail to teach or suggest, alone or in combination, the claimed invention, as required under Lea and MPEP § 2131. For at least these reasons, the PTO's rejections of claim 1 and all claims of Group IV should be reversed.

Claims Group V:

Referring to the fifth group of claims (claims 18-21, 25-29), these claims have been rejected under 35 U.S.C. § 102(e) as being unpatentable over U.S. Patent No. 6,115,373 to Lea. For purposes of this discussion, claim 18 of Group V will be addressed.

Appellants respectfully maintain that Lea fails to teach or suggest a low latency switched network including a first switch, the first switch being a buffer-less switch coupling the plurality of sending and receiving nodes, <u>the buffer-less switch having a fixed forwarding delay</u> for all packets sent from one of the sending nodes and successfully received by one of the receiving nodes, as recited by claim 18. The Final Office Action states that Lea "mentions nothing about changing the forwarding rates of the incoming packets, so the system forwards the incoming packets at the same rate that they were received." The Final Office Action states that

> if the system of Lea receives packets at a constant bit rate, then the system will be forwarding the incoming packets at the same constant bit rate, which is also known as a fixed rate. Lea discloses in col. 4, lines 1-15, that the system can handle real-time traffic. If the incoming packets are real-time traffic, then they will be forwarded at this same fixed rate as well. Therefore Lea discloses a system that can have a fixed forwarding rate.

Applicants respectfully disagree. A "real-time system" is a "system in which transactions are processed as they occur." RUDOLPH F. GRAF, MODERN DICTIONARY OF ELECTRONICS, at 628 (7th ed. 1999).

The Final Office Action fails to point out how a system that processes data as it is received teaches or suggests that the switch of Lea has a fixed forwarding rate. "A claim is anticipated only if each and every element as set forth in the claim is found, either <u>expressly or inherently described</u>, in a single prior art reference. *Verdegaal Bros. v. Union Oil Co. of*

California, 814 F.2d 628, 631, 2 USPQ2d 1051, 1053 (Fed. Cir. 1987)." See MPEP § 2131 (emphasis added). Lea fails to teach or suggest, expressly or inherently, forwarding all packets that are successfully delivered through output ports of the buffer-less switch to the receiving nodes, through the buffer-less switch with a fixed forwarding rate, as recited by claim 18.

Thus, the references fail to teach or suggest, alone or in combination, the claimed invention. For at least these reasons, the PTO's rejections of claim 18 and all claims in Group V should be reversed.

CONCLUSION

For at least the foregoing reasons, Appellants' presently claimed invention would not have been anticipated under 35 U.S.C. § 102(a) or obvious under 35 U.S.C. § 103(a) as set forth in the Final Office Action. Accordingly, this Board is respectfully requested to reverse the rejection of claims 1-30 and direct this application to be issued.

CERTIFICATE OF MAILING OR TRANSMISSION

I hereby certify that, on the date shown below, this correspondence is being

☒ deposited with the US Postal Service with sufficient postage as first class mail, in an envelope addressed to Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

☐ facsimile transmitted to the US Patent and Trademark Office.

Mark Zagorin — 8/11/04

Mark Zagorin — Date

EXPRESS MAIL LABEL: ______

Respectfully submitted,

Mark Zagorin

Mark Zagorin, Reg. No. 36,067
Attorney for Appellant(s)
(512) 338-6311
(512) 338-6301 (fax)

APPENDIX OF CLAIMS INVOLVED IN THE APPEAL

1. (Original) A method for communicating packets between one of a plurality of sending nodes and one of a plurality of receiving nodes of a switched network, the switched network including a buffer-less switch coupling the sending nodes and the receiving nodes, the method comprising:

   transmitting packets from respective sending nodes to respective input ports of the buffer-less switch; and

   forwarding all packets that are successfully delivered through output ports of the buffer-less switch to the receiving nodes, through the buffer-less switch with a fixed forwarding rate.

2. (Original) The method as recited in claim 1 wherein each receiving node sends an acknowledge to a respective sending node at a predetermined time with respect to sending a corresponding packet, to indicate successful delivery of the corresponding packet to the sending node, thereby providing a fixed time for the sending node to know whether a packet was successfully transmitted.

3. (Original) The method as recited in claim 2 wherein the sending node determines that transmission of a packet was unsuccessful by checking if the acknowledge was returned after the predetermined time has elapsed.

4. (Original) The method as recited in claim 3 wherein the receiving node sends a no acknowledge (nack) at the predetermined time to the sending node on detection of an error condition in receipt of the packet.

5. (Original) The method as recited in claim 4 wherein the error condition detected by the receiving node is one of a buffer overflow and a checksum error.

6. (Original) The method as recited in claim 2 wherein unsuccessful transmission is determined by a timeout indicating that an acknowledge failed to arrive after the predetermined time has elapsed.

7. (Original) The method as recited in claim 1 further comprising:

for each packet being sent over the switched network, requesting respective transmission paths through the switched network to one of the receiving nodes;

allocating one of the transmission paths to a first requester with respect to arrival time in the buffer-less switch, the first requester requesting the one transmission path, and ignoring any other requests for the one transmission path until the one transmission path again becomes available; and

if multiple requests collide by requesting a switch resource simultaneously, selecting a first packet associated with one of the requests as a winning packet and dropping any other packets associated with requests other than the one request.

8. (Original) The method as recited in claim 7 wherein the requests for transmission paths are contained within the packets sent into the network and extracted after entry into the switch.

9. (Original) The method as recited in claim 7 further comprising selecting the winning packet according to at least one of a random basis and a round robin basis.

10. (Original) The method as recited in claim 7 further comprising selecting the winning packet according to a fairness criteria having an objective to allocate to each input port an equal share of bandwidth at each output port.

11. (Original) The method as recited in claim 7, wherein no buffer space is allocated in a receiving node before a packet is sent, thereby simplifying switch overhead.

12. (Original) The method as recited in claim 11 wherein if the receiving node detects a buffer overflow, the receiving node sends a no acknowledge packet (nack) to the sending node indicating that a packet associated with the buffer overflow was not successfully received.

13. (Original) The method as recited in claim 1 wherein communicating packets between one of the sending nodes and one of the receiving nodes further comprises:

the sending node writing a packet to a send register;

the sending node polling a status register to determine if transmission of the packet was successful; and

if transmission was unsuccessful rewriting the packet to the send register.

14. (Original) A method for utilizing a switch of a switched network comprising:

forwarding packets at a fixed rate on a first come first served basis from respective input ports through the switch to respective output ports; and

if a first and second packet simultaneously request a switch resource, selecting one of the first and second packets a winner and one a loser, the winner obtaining the switch resource, and dropping the loser.

15. (Original) The method as recited in claim 14 wherein the selecting is determined according to at least one of a random selection and a round robin selection.

16. (Original) The method as recited in claim 14 wherein low latency packets are transmitted on the switch.

17. (Original) The method as recited in claim 14 wherein a sending node can positively determine after a fixed delay with respect to sending of the packet that a packet was successfully transmitted across the switched network.

18. (Original) A computing system comprising:

a plurality of sending and receiving nodes;

a low latency switched network including a first switch, the first switch being a buffer-less switch coupling the plurality of sending and receiving nodes, the buffer-less switch having a fixed forwarding delay for all packets sent from one of the sending nodes and successfully received by one of the receiving nodes.

19. (Original) The computing system as recited in claim 18 further comprising a second switched network, including a second switch coupled to the plurality of sending and receiving nodes.

20. (Previously presented) The computing system as recited in claim 19 wherein the first switch carries scheduling information for a storage device and the second switch carries bulk traffic for at least one of storage and retrieval on the storage device.

21. (Original) The computing system as recited in claim 18 wherein each sending node includes a send register written into by a sending node to send data across the network.

22. (Original) The computing system as recited in claim 21 wherein each sending node includes a status register indicating whether a transfer across the network completed successfully.

23. (Original) The computing system as recited in claim 22 wherein the status register includes a field indicating a type of failure.

24. (Original) The computing system as recited in claim 22 wherein a sending node rewrites data into the send register if a transfer across the network for the data completed unsuccessfully.

25. (Original) The computing system as recited in claim 18 wherein the buffer-less switch further comprises:

a plurality of input registers coupled to respective input ports;

switch control logic, coupled to the input registers and responsive to packet information stored in the registers, to allocate output ports on the switch according to the packet information;

and wherein the switch control logic is responsive to allocate output ports on a first come first served basis.

26. (Original) The computing system as recited in claim 25 wherein respective packet information provided to the switch control logic constitutes respective requests for output ports,

and if a first and second request for an output port path collide by requesting the output port at the same time, the switch control logic responds by selecting one of the requests as a winner and dropping a packet associated with the second request.

27. (Original) The computing system as recited in claim 25 further comprising output registers in the buffer-less switch coupled to receive data selected by respective selector circuits selectively coupled to respective ones of the input ports.

28. (Original) The computing system as recited in claim 26 wherein the switch control logic selects the winner according to at least one of a random basis and a round robin basis.

29. (Original) The computing system as recited in claim 18 wherein the low latency switched network includes a plurality of cascaded buffer-less switches, thereby forming a multi-stage buffer-less switch.

30. (Previously presented) A switched network comprising:

means for forwarding packets at a fixed rate on a first come first served basis from respective input ports through a switch to respective output ports; and

means for selecting one of a first and second packet as a winner and one a loser, the winner obtaining a switch resource, and dropping the loser, if the first and second packet simultaneously request the switch resource.